REPORT DOCUMENTATION PAGE

AFRL-SR-AR-TR-05-

0407

1. AGENCY USE ONLY *(Leave blank)*	2. REPORT DATE	3. REPORT ... FINAL REPORT

4. TITLE AND SUBTITLE
SIMULATION AND OPTIMIZATION METHODOLOGIES FOR MILITARY TRANSPORTATION NETWORK ROUTING AND SCHUDULING AND FOR MILITARY....

5. FUNDING NUMBERS
F49620-02-1-0099

6. AUTHOR(S)
DR. ERVIN Y. RODIN

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
CENTER FOR OPTIMIZATION AND SEMANTIC CONTROL
DEPARTMENT OF SYSTEMS SCIENCE AND MATHEMATICS
CAMPUS BOX 1040, WASHINGTON UNIVERSITY
ONE BROOKINGS DRIVE
ST. LOUIS, MO 63130-4899

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
AIR FORCE OFFICE OF SCIENTIFIC RESEARCH
801 NORTH RQANDOLPH STREET
ARLINGTON, VA 22203-1977

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION AVAILABILITY STATEMENT
Distribution Statement A. Appoved for public release; distribution is unlimited.

12b. DISTRIBUTION CODE

13. ABSTRACT *(Maximum 200 words)*
We are commencing this final report by describing briefly the Attachment to the report. The attachment itself is a relatively short description of what we consider one of our major accomplishments during this research: the unification of simulation and optimization for airlift networks.

14. SUBJECT TERMS	15. NUMBER OF PAGES
	16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT

REPORT DOCUMENTATION PAGE	*Form Approved OMB No. 0704-0188*

1. AGENCY USE ONLY *(Leave blank)*	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED FINAL REPORT

4. TITLE AND SUBTITLE
SIMULATION AND OPTIMIZATION METHODOLOGIES FOR MILITARY TRANSPORTATION NETWORK ROUTING AND SCHUDULING AND FOR MILITARY....

5. FUNDING NUMBERS
F49620-02-1-0099

6. AUTHOR(S)
DR. ERVIN Y. RODIN

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
CENTER FOR OPTIMIZATION AND SEMANTIC CONTROL
DEPARTMENT OF SYSTEMS SCIENCE AND MATHEMATICS
CAMPUS BOX 1040, WASHINGTON UNIVERSITY
ONE BROOKINGS DRIVE
ST. LOUIS, MO 63130-4899

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
AIR FORCE OFFICE OF SCIENTIFIC RESEARCH
801 NORTH RQANDOLPH STREET
ARLINGTON, VA 22203-1977

10. SPONSORING/MONITORING AGENCY REPORT NUMBER

11. SUPPLEMENTARY NOTES

12a. DISTRIBUTION AVAILABILITY STATEMENT
Distribution Statement A. Appoved for public release; distribution is unlimited.

12b. DISTRIBUTION CODE

13. ABSTRACT *(Maximum 200 words)*
We are commencing this final report by describing briefly the Attachment to the report. The attachment itself is a relatively short description of what we consider one of our major accomplishments during this research: the unification of simulation and optimization for airlift networks.

14. SUBJECT TERMS	15. NUMBER OF PAGES
	16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT

FINAL REPORT

Submitted to
Air Force Office of Scientific Research
801 North Randolph Street
Arlington VA 22203-1977

by

Ervin Y. Rodin, P.I.
Professor and Director

Center for Optimization and Semantic Control
Department of Systems Science and Mathematics
Campus Box 1040, Washington University
One Brookings Drive
St. Louis, MO 63130-4899

in connection with
Grant AFOSR F49620-02-1-0099

SIMULATION AND OPTIMIZATION METHODOLOGIES FOR MILITARY TRANSPORTATION NETWORK ROUTING AND SCHEDULING AND FOR MILITARY MEDICAL SERVICES

August 3, 2005

1. INTRODUCTION

We are commencing this final report by describing briefly the Attachment to the report. The attachment itself is a relatively short description of what we consider one of our major accomplishments during this research: the unification of simulation and optimization for airlift networks.

Specifically, our main contribution in this regard is a novel method of introducing optimization into military transportation analysis through the embedding of optimization techniques into simulations to an extent that has never been done before. These optimization techniques replace the simple rule-based decision-making strategies commonly used in military simulations.

Our approach is evolutionary in nature, not revolutionary. We begin with a tool that transportation analysts are comfortable with, namely a simulation. Then we embed optimization strategies into the simulation by making calls to an appropriate optimization model each time a decision needs to be made (i.e. which route an aircraft should take). An optimal decision is made based on some objective function and passed back to the simulation.

With this approach, analysts would see very little change in the operations of the simulations. We do not require repeated iterations of the simulation each time the data set changes to find "optimal" values of parameters used in the simulation. Nor do we try to model extremely complex, stochastic systems with a closed set of equations.

We applied our approach to the Airlift Network Problem from the United States Air Force's Air Mobility Command (AMC) where the basic objective is to simulate the delivery of a list of cargo with a given fleet of aircraft while minimizing the amount of late cargo. We then demonstrate how decisions are modeled using linear and non-linear models.

We also compare results from our optimization strategies to results using simple rule-based decision-making strategies. Through this we demonstrate how our optimization approach results in more cargo being delivered on-time, thus these strategies produce more desirable results with the same amount of resources. We hope that our results will convince analysts to consider using optimization strategies in their analysis.

2. OBJECTIVES.

The purpose of this present research was to develop a generic model and methodology for analyzing and optimizing large-scale air transportation networks, including both their routing and their scheduling. We proposed to achieve this aim in part by studying several specific examples of current problems of this type, arising in the operations of the Air Mobility Command (AMC) at Scott AFB; and in part by developing further the various paradigms that we had employed successfully in the past in similar contexts. These include the utilization of the classical mathematical methodologies of Linear and Integer Programming, ... time dependent integer programming.... We continued to collaborate with scientists from Scott AFB; indeed, our main attempt was to improve various aspects of AMC's Mobility Analysis Support System (MASS). We intend for these aspects to serve as the particular paradigms for the general model and methodology to be developed. In addition, we undertook to study the efficiency of certain aspects of the military medical support services, through simulations, attempts at optimization, and the comparative evaluation of various medical services provided by the military.

3. STATUS OF EFFORT.

Quote from our Progress Report, dated September1, 2002:

One of our major projects, a doctoral dissertation on the subject of

Intelligent Transportation Scheduling: Heuristic and Sequential Optimization of Simulated Transportation Systems

has been completed. Its principal application was to the

Strategic Brigade Airdrop Operation.

The results were presented to AMC/XPY and to US TRANSCOM at Scott AFB, and it was received favorably. Three copies of the dissertation were also sent to AFOSR.

We are also continuing work on a novel methodology, applied both to military transportation systems and to some civilian projects, through which we are attempting to incorporate optimization into simulations in a novel way.

Finally, we are continuing our joint effort with the Medical Defense Partnership for Reinvention (MDPR). One graduate and three undergraduate students spent an entire summer at the Internal Medicine Clinic of the Scott AFB hospital, and developed a simulation of their activities. These results were presented to various entities at the hospital and at MDPR. The project was received very well by all, and we were encouraged to develop a generalization of it, applicable to all DOD medical facilities. In addition, we also developed a methodology to evaluate the relative efficiency of a large number of DOD clinics and hospitals. Our methodology allows the users to pinpoint problem areas and ways to introduce corrections.

The results of this medical project will be published. To accompany these results, we also developed an Excel add-in, called the DEA Solver.

Quote from our Progress Report, dated September 1, 2003:

During this reporting period we succeeded in our attempts to incorporate optimization into simulation in a novel and seamless way. Furthermore, we are able to exercise this methodology by using off the shelf software: namely, PROMODEL for simulation, ILOG for optimization and EXCEL for data collection and organization.

We are providing, as attachments, two reports on this project: one is an application to military transportation and the other to a civilian chemical plant.

Also, we are continuing our joint effort with the USAF Command Surgeon's project. This past summer two graduate and one undergraduate student spent an entire period at the Emergency Room of the Scott AFB hospital, and developed a simulation of their activities. These results were presented to various entities at the hospital. The project was received very well by all, and we were encouraged to develop a generalization of it, applicable to all DOD medical facilities, and to further generalize it to the much larger civilian Emergency Rooms.

Quote from our Progress Report, dated November 9, 2004:

During this reporting period we succeeded in our attempts to incorporate optimization into simulation in a novel and seamless way. Furthermore, we are able to exercise this methodology by using off the shelf software: namely, PROMODEL for simulation, ILOG for optimization and EXCEL for data collection and organization.

In particular, the novel way in which we constructed our routing and scheduling algorithm can be briefly described as follows:

- *Relevant data (i.e., TPFDD's) are entered into the simulation;*

- *The simulation (using PROMODEL) begins to run;*
- *When a choice is available for action, the simulation calls an optimization algorithm (in ILOG);*
- *The results of the optimization are returned to the simulation, which continues until the next choice becomes available;*
- *Repeat this process until the simulation is completed.*

There are several advantages to our methodology. Two of the most important ones are these:

- *The methodology is completely transparent to the user, in the sense that all he has to do is run a simulation as usual. This means that since currently routing and scheduling is done purely by (a non-optimized) simulation, users need not be concerned with, nor even become familiar with the underlying optimization mechanism;*
- *We have shown that considerable time and monetary savings result from employing our methodology.*
-

Several months ago we provided three copies of a doctoral dissertation by Brian Albright, entitled

An Embedded Optimization-Simulation Approach to Dynamic Pickup and Delivery Problems

This dissertation contains a complete description of the above project.

In addition, we are now in the process of extending these results even further, by considering the stochasticity of many elements of routing and scheduling, with the aim of incorporating these considerations in our software.

Finally, we are continuing our effort with the USAF Command Surgeon's project. In particular, having completed our projects at the military medical facilities, it became necessary to compare these results to their civilian counterparts. Therefore, this past summer one graduate and one undergraduate student spent their entire three months at the Emergency Room of the Missouri Baptist Hospital, and developed a simulation of their activities. These results were presented to various entities at the hospital. The project was received very well by all, and we were encouraged to develop a generalization of it, applicable to all DOD medical facilities, and to further generalize it to even larger civilian Emergency Rooms.

4. ACCOMPLISHMENTS/NEW FINDINGS

See point 3. above

5. PERSONNEL ASSOCIATED WITH THIS RESEARCH

Faculty:

Professor Ervin Y. Rodin (PI)

Graduate Students:

T. Eugene Day
Gregory Grindey
Brian Albright
Yong Huang
Ashoka Polpitiya
Xiaohu Jin

Undergraduate Students:

Changjae Lee
Paulo Pirondi
Shirley Birman
Daniel Livengood
Haruka Kakimoto

6. PUBLICATIONS

None so far.

7. INTERACTIONS/TRANSITIONS

Joint development with AMC/XPY, US TRANSCOM, the Medical Defense Partnership for Reinvention (MDPR), Tyco Healthcare/Mallinckrodt, Scott AFB Hospital, Duke University Medical School, Missouri Baptist Hospital.

8. NEW DISCOVERIES, INVENTIONS OR PATENT DISCLOSURES

As described in 3. above; no patents.

9. HONORS/AWARDS

None

10. ATTACHMENTS

We attached to our previous Progress Reports items **1-3** below; and are attaching now #**4**.

1. 3 copies of the dissertation

   Intelligent Transportation Scheduling: Heuristic and Sequential Optimization of Simulated Transportation Systems

   were provided to the AFOSR and to Scott AFB;

2. 3 copies of the dissertation

   An Embedded Optimization-Simulation Approach to Dynamic Pickup and Delivery Problems

   were provided to the AFOSR and to Scott AFB;

3. 3 copies of the paper

   Application of DEA to Medical Clinics,

   and

   A Guide to Using DEA Solver

   were also provided to the AFOSR and to MDPR.

4. Finally, we are also enclosing here an abstracted version of one of the main aspects of the dissertation in point **2** above, as mentioned in our Introduction, with the title

   Simulation and Optimization of an Airlift Network.

Attachment 4:

SIMULATION AND OPTIMIZATION OF AN AIRLIFT NETWORK

INTRODUCTION

Transportation analysts at Air Mobility Command (AMC) and U.S. Transportation Command (USTC) are often asked questions such as, "how much cargo can be delivered with the available resources," or, "how many aircraft are needed to meet the demand?" For years, the tool of choice to answer questions such as these for analysts at both AMC and USTC has been simulation. Simulations allow analysts to model extremely complicated dynamic systems and use rather simple heuristic rules to assign cargo to vehicles, choose routes, etc. These simulations yield detailed information on the activities of ships, aircraft, and other vehicles and the activities of airports and seaports that can be analyzed during and after the simulation.

For years, researchers have tried to get the military to implement optimization strategies in their analysis. Several researchers have demonstrated how optimization techniques can be used in military transportation analysis. One of the most notable examples is the THRUPUT model developed by Morton, Rosenthal, and Weng (Morton, 1996). THRUPUT is a linear model designed to determine the maximum on-time throughput of an airlift network with a given fleet of aircraft. It uses pure integer modeling techniques and search engines with no simulation.

Another optimization model similar to THRUPUT is the NPS/RAND Mobility Optimizer (NRMO) used by AMC (Rink, 1998). NRMO is a large-scale linear program written in the modeling language GAMS (General Algebraic Modeling System) that models airlift networks. Its objective is to maximize the delivery of cargo by minimizing late and non-delivered cargo.

The main advantages of models such as these are that they use sophisticated search strategies and modeling tools (i.e. CPLEX). The disadvantages include the fact that many simplifications must be made when modeling, resulting in a low-fidelity model. Also, the model returns limited information on the specific activities of individual aircraft, which hinders the analysis of the solution.

Optimization tools such as THRUPUT and NRMO have failed to become a major part of the analysis tools used by AMC and USTC. One reason for this is that the problems studied by AMC and USTC are far too large and complex to be modeled with a high-fidelity explicit mathematical model. Another reason, and certainly the most powerful reason to overcome, is a resistance to change. Simulations have been the only tool used so far and people simply don't want to change.

Therefore, any attempt at incorporating optimization into the set of analysis tools must be *evolutionary*, not *revolutionary*. If analysts see only a gradual change to incorporating optimization, they will not be as resistant to change.

In our solution of the Airlift Network Problem, we present a truly novel, evolutionary approach. Our solution begins with a tool analysts are comfortable with, namely a simulation, and then we embed very powerful and sophisticated optimization techniques into the simulation. We call the approach the "integration of simulation and optimization."

Through this approach we will illustrate two points. First we will show that these very complex problems can be modeled with high fidelity using available optimization tools. Second we will show that the use of optimization strategies will produce more desirable results than the use of pure simulation (i.e. more cargo can be delivered with the available resources, or fewer aircraft are needed to meet demands).

Others, such as Powell (Powell, 2001), have also used optimization in conjunction with simulation for the purposes of modeling military transportation networks. However, our approach differs considerably from that of Powell.

One fundamental deficiency with virtually all simulations is the use of simple decision-making strategies such as if..then..else logic or look-down list strategies. Simulations used by the Air Force often use very simple look-down list strategies for making rather

complicated decisions that can have a great affect on the system as a whole. This leads to options being ignored, some of which may be better than the ones chosen, and in general, poor decision-making abilities.

Powell's approach can be seen more as an *iteration* of simulation and optimization rather than an *integration*, or embedding. Powell retains the use of common decision-making strategies in simulations. Through repeated iterations of the simulation, artificial intelligence methods are used to "learn" the best strategies for making decisions, or the best values of parameters used in making decisions. He calls this tool "The Optimizing-Simulator." This approach falls into the general category of "simulation optimization."

One fundamental problem with this type of approach is that when the data set changes, the parameters used in making decisions will most likely also change. Therefore a lengthy process of running the iterations again must be undertaken. Another problem is the fact that this approach is significantly different than what is currently being used by the Air Force. Air Force analysts are not used to running several iterations of their simulation to find new values of parameters each time the data set changes.

Our general approach is to replace the simple decision-making strategies in simulations with the use of more sophisticated mathematical modeling and optimization techniques. Each time a decision needs to be made during the simulation, the simulation software will make a call to the optimization software, which will then make the appropriate (locally) optimal decision and pass the decision back to the simulation.

This approach has the distinct advantage of being data independent. When the data changes, no changes need to be made to the models. Also, the analyst will see very little change in the actual operation of the simulation. No repeated iterations are necessary. The basic goal of this approach is to produce more desirable results from the simulation by improving its decision-making ability.

Grindey and Cusick have also embedded optimization into simulations, although to a somewhat lesser extent than we have. In his doctoral thesis, Travis Cusick (Cusick, 2000) models and optimizes a military airfield system with a tool named the Base Resource And Capability Estimator (BRACE) (Note: BRACE has been further developed by USTC and is now the Airport Simulation Tool (AST) component of the modeling and simulation software TRANS-PORT). He decomposes the system into several subsystems including aircraft arrivals, fuel trucks, and parking spots. He then develops strategies to optimize each subsystem.

Greg Grindey in his doctoral thesis (Grindey, 2002) simulates a military brigade airdrop. In the simulation, different aircraft may perform different roles. Grindey uses an integer model to optimize the assignment of roles and compares the results to heuristics currently used by AMC.

THE AIRLIFT NETWORK PROBLEM

The basic objective of the Airlift Network Problem from AMC is to simulate the delivery of a list of cargo with a given fleet of aircraft while minimizing the amount of late cargo. Included in the simulation are stochastics representing variability in ground times and simple aircraft breakdowns.

The list of cargo is comprised of the level 2 detail of a Time Phased Force Deployment Data (TPFDD) document. For each record, the important parameters include total liquid tons of Outsized, Oversized, and Bulk cargo; the number of passengers; Aerial Port of Embarkation (APOE); Aerial Port of Debarkation (APOD); Available to Load Day (*ALD*); Required Delivery Day (*RDD*); and a Cargo Commodity Code. The APOE and APOD are the airbases where the cargo is to be picked up and delivered, respectively. The *ALD* is the earliest simulation day the cargo can be loaded. The *RDD* is the latest simulation day the cargo can be delivered. The *ALD* and *RDD* are taken as hard and soft constraints, respectively. The Cargo Commodity Code is a rough measurement of the

volume of the cargo and is used in conjunction with the Payload Target Data File to determine how cargo is loaded on individual aircraft.

The fleet of aircraft consists of 10 each of Wide Body Passenger (WBP), Wide Body Cargo (WBC), C-17, and C-5 aircraft. WBP and WBC are both civilian Boeing 747 aircraft that can carry exclusively passengers and bulk cargo, respectively. C-17 and C-5 aircraft specialize in carrying Outsized and Oversized cargo, but can also carry Bulk cargo and passengers.

Potential routes between different airbases are predefined. The world is broken up into different regions denoted by integers. For each possible source region (a region containing the starting point of a route) and sink region (a region containing the destination) and each aircraft type, a list of possible routes composed of a sequence of waypoints (imaginary turning points in the sky) and en routes (refueling locations) is given.

THE SOLUTION

Our solution is an instantiation of the concept of integrating simulation and optimization. The basic movement of aircraft and loading and unloading of cargo, including stochastics, are modeled in a discrete-event simulation written using the commercial software product ProModel. Each time a decision needs to be made, the simulation passes appropriate data to the optimal search engine ILOG where the decision to be made is modeled, and an optimal decision is made based on an appropriate objective function and passed back to the simulation.

Two main types of decisions are made during the simulation: assignment of aircraft to cargo and the selection of routes. A binary, linear assignment model solved with CPLEX MIP is used to make the first type of decision, while a non-linear scheduling model solved with ILOG Solver is used to make the second type.

THE SIMULATION

A typical simulation will have "resources" and "entities." Resources are objects that perform actions on entities (i.e. vehicles pick up customers, machines process raw material, etc.). Aircraft in this simulation are modeled as entities and there are no resources. Cargo exists only as numbers in arrays and is not modeled as entities. This is a somewhat radical way of modeling the system. Common sense says to model the cargo as entities and aircraft as resources. Several factors caused us to choose this type of model including:

- In a typical simulation entities "call" resources and the resources move to the entities, so the movement of resources is controlled directly by the entities. In this model, ILOG is used to assign cargo to aircraft, so the cargo will not directly "call" the aircraft.
- Once an aircraft on-loads cargo its movement to an APOD and the end of the trip is controlled by a schedule created in ILOG, not by the cargo.
- In ProModel, entities can be moved from location to location using processing logic. Resources can move only after being called by an entity.
- Although the objective is to minimize the amount of late cargo, the movement of aircraft is really what is important since cargo can't be delivered without the movement of aircraft.

The basic path of aircraft in the simulation is shown in Figure 1. Each aircraft begins the simulation at its designated Home Base. Once it receives a cargo pickup assignment, it flies to the appropriate APOE, loads the cargo, and then flies to the appropriate APOD and off-loads the cargo. The aircraft will then fly to one of its designated Recovery Bases. In between each type of base, the aircraft may stop at anywhere from zero to three en routes for fuel.

Once it finishes at the Recovery Base, a random draw is made to determine if the aircraft requires maintenance at its Home Base. Ten percent of the time and aircraft will require

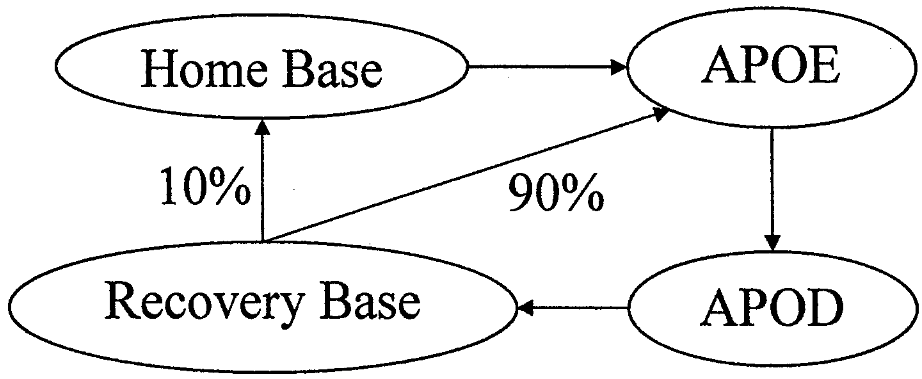

Figure 1: Aircraft Paths

maintenance and the aircraft will fly to its Home Base for maintenance, which is modeled as requiring 10 hours. The remainder of the time, the aircraft will get its next cargo pickup assignment and fly to the appropriate APOE. If the aircraft does not have a next assignment, it will fly to its Home Base and wait for an assignment.

When an aircraft is ready to leave a Home Base or APOE to start a "trip," it passes control to ILOG which selects a route and creates a "schedule" for that trip. ILOG also calculates flight times and cargo capacities. The schedule contains the sequence of locations at which the aircraft is to stop and the approximate times it will arrive and leave each location, as shown in Figure 2. This schedule does not tell the simulation exactly when the aircraft is to arrive and leave each location. The simulation determines these exact times.

Activity	Base	Expected Start	Expected End	Scheduled Start	Scheduled End	Fuel
Load at APOE	KDOV	0	195	0	195	0
Stop at En route	ETAR	687	882	687	882	2996
Off-load at APOD	OBBS	1298	1553	1298	1553	2563
Recovery Base	LEMO	1755	1950	1755	1950	1962
Return to APOE	KDOV	2412	2412	2412	2412	2838

Figure 2: A Typical Schedule

Individual activities performed at each stop are not modeled (i.e. landing, parking, loading cargo, etc.). Rather, the total time spent on the ground is modeled with a simple "wait" statement. This command instructs the aircraft to wait at the base for a certain amount of time and then immediately leave. The actual amount of time an aircraft spends performing an activity on the ground is a random variable defined by:

(1) $$\text{Wait Time } = \text{GT}(0.95 + 0.1 * X)$$

where X is a random variable with a lognormal distribution and $E[X] = \sigma(X) = 1$ and GT is the standard planning factor Ground Time for that activity and aircraft type.

Wait Time is defined in this way because the planning factor ground time (GT) is located at the peak of graph of the density function as shown in Figure 3. There is a probability that the wait time is less than the GT, although it is rather small. There is a much higher probability that the wait time is higher than the GT. The average wait time it 1.05 * GT, or the GT plus 5%.

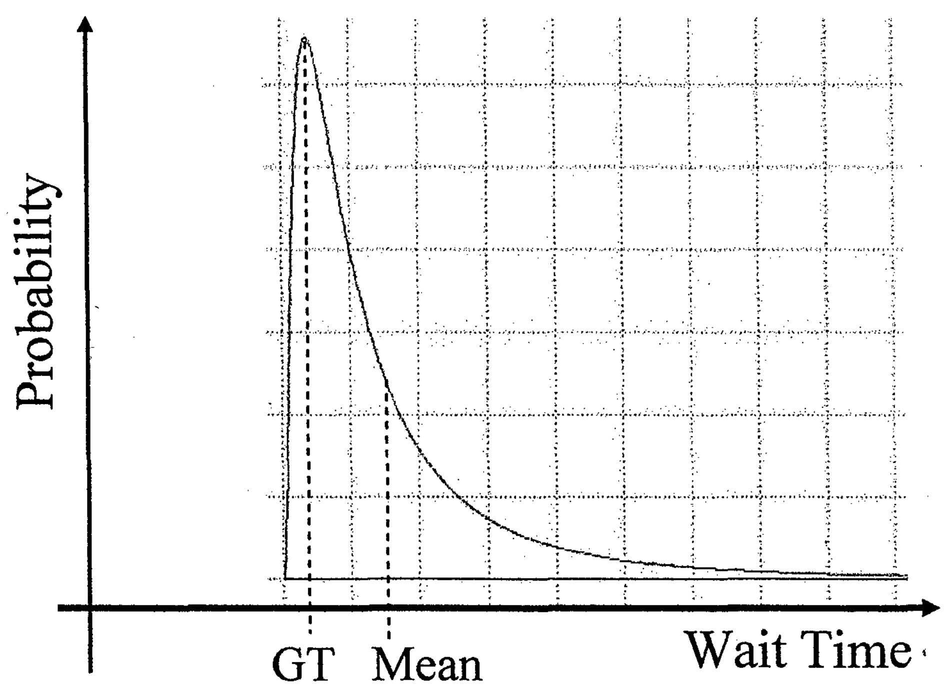

Figure 3: Wait Time Density Function

When an aircraft gets its schedule, it begins performing the first activity at the time indicated in the Expected Start column. As soon as it is finished at the first base, it "updates its schedule" by adding (Current Simulation Time – Expected End Time) to the expected end time of the current activity and to the expected start and end times of all future activities. It then flies to the next base in the schedule and updates its schedule in a similar manner as soon as it lands. The times in the scheduled column are never changed; these numbers are used for analysis later. The purpose of "updating a schedule" is to get better estimates of where each aircraft is at any point in the simulation.

Waypoints are not modeled in the simulation due to limitations with the simulation software so aircraft always fly in straight paths between locations. However, waypoints are taken into account when calculating distances in ILOG, which are used to calculate flight times when routes are chosen, so the flight times in the simulation are accurate.

Fuel is not modeled in the simulation either, but it is taken into consideration when choosing routes and creating schedules. Each base has a limited daily fuel dispensing capacity, but these capacities are used only as planning factors and do not represent any physical constraints.

Parking and MOG

Maximum On Ground (MOG) is a simple description of the capacities of an airbase determined by factors such as parking space and the number of available resources, such as Material Handling Equipment (K-loaders) and fuel trucks. It also takes into account factors such as the type of fuel pumping resources (i.e. fuel trucks or in-ground pumps) and the physical distances between fuel pumps.

MOG can be interpreted to mean the maximum number of aircraft that can be serviced simultaneously within their allotted ground time. Each aircraft type has an allotted ground time for each type of activity: on-load, en route, and off-load. In reality, an

airbase could exceed its MOG constraint, but then it would not be able to service all the aircraft in their allotted ground time.

One other factor that MOG takes into consideration is the size of the aircraft. Aircraft types are divided into two general body-type categories: Narrow Body or Wide Body. In our simulation, all aircraft types except C-17 are considered Wide Body.

MOG consists of four numbers and a logic type. The numbers are divided into two categories: Parking MOG and Working MOG. Each category is further divided into two sub-categories: Narrow Body MOG and Working Body MOG. These four numbers are designated: Narrow Body Parking (NBP), Wide Body Parking (WBP), Narrow Body Working (NBW), and Wide Body Working (WBW). Parking MOG describes the number of each body-type that could be parked at the airbase at any time. Working MOG describes the number of each body-type that could be serviced at any time and is a subset of parking MOG, i.e. any aircraft that is being serviced is also parked. Therefore, parking MOG is always at least as large as working MOG.

The logic is given as "AND" or "OR." To illustrate the "OR" logic take, for example, a parking MOG of "5 OR 3." This means that there could be a maximum of 5 narrow body aircraft or 3 wide body aircraft, or some combination thereof, parked at a time. This is modeled as meaning that a narrow body aircraft takes 1/5 of the total parking capacity and a wide body aircraft takes 1/3 of the capacity with the total capacity used never exceeding 1.

A parking MOG of "5 AND 3" means that no more than 3 wide body aircraft can be parked at a time and the total number of narrow body and wide body aircraft cannot exceed 8. With the "AND" logic, parking spots can be thought of as being divided into two categories, wide body and narrow body. Only a narrow body aircraft can fit into a narrow body parking spot, but narrow body aircraft can also fit into a wide body parking spot.

In the simulation, both parking and working MOG are considered. When an aircraft arrives at an airbase, it will wait in the air until an appropriate amount of parking MOG becomes available. As soon as it becomes available, the aircraft will land and then wait until an appropriate amount of working MOG becomes available. As soon as it becomes available, the Wait Time that models the activities performed at the base begin. As soon as this time has passed, the aircraft leaves for the next airbase in its schedule even if it is ahead of, or behind, schedule. At that point the parking and working MOG are released for use by another aircraft.

CHOOSING ROUTES AND CREATING SCHEDULES

Routes are modeled using the software ILOG. ILOG is a suite of modeling and optimization components including, CPLEX, Solver, Scheduler, and OPL (Optimization Programming Language) Studio. CPLEX is the well-known linear search engine. Solver is a non-linear, integer, constraint-programming search engine. Scheduler is a set of constructs for modeling scheduling problems. The models created with Scheduler are non-linear integer models solved with Solver. OPL Studio is ILOG's own language and compiler. All optimization models in this paper were written with OPL Studio.

Routes are modeled using constructs from Scheduler. The two main concepts behind Scheduler are Activities and Resources. Activities are objects composed of a start time, end time, and duration with the constraint that the end time equals the start time plus the duration. Each one of these quantities must be an integer and could be explicitly defined or left as a variable and have constraints placed on it. Resources are entities required for the execution of an activity. Resources must have integer capacities and be used by activities in integer quantities.

Activities can "require" or "consume" resources. Requiring a resources means that a sufficient quantity of the resource must be available over the duration of the activity. Consuming a resource means that a sufficient quantity must be available from the start of the activity until the schedule horizon, which is defined by the user.

ILOG also has a very powerful function named maxCapacity that allows the modeling of varying resource capacities. For instance, to model the instance where a resource named "Fuel" has a maximum capacity of 500,000, but 100,000 of this is used from time 250 to time 500, we would use the line:

maxCapacity(Fuel,250,500,100000);

Each aircraft carries a schedule with it at all times. Several different types of schedules can be created during the simulation, depending on the circumstances of the aircraft. The typical schedule is one where the aircraft begins at an APOE where cargo is loaded, stops at an APOD where cargo is off-loaded, stops at a Recovery Base, and then returns to either another APOD or a Home Base. This is called a "Regular" schedule. We will illustrate the process of creating a Regular schedule.

The first step in the process of choosing a route and creating a schedule is to get a measurement of how much, and when, the parking and fuel resources of each base are going to be used in the near future. To do this, a subroutine in the simulation looks at the expected start and end times in the schedules of all aircraft to count the number of narrow body and wide body aircraft expected to be at each base in each 15-minute block of time for two days into the future. These data are stored in an array named "Expected Occupants". Similarly, it adds the total amount of fuel expected to be consumed at each base in each simulation day for four days into the future and stores it in an array named "Expected Fuel Usage." These two arrays are written to a flat-text data file in a format that can be read by ILOG.

The next step is to write a flat-text data file containing the aircraft type, APOE, APOD, and destination (another APOE or Home Base). Control is then passed to ILOG, which first looks through a data file to find a list of possible routes and recovery bases. Then, the shortest possible schedule for each possible route is created. The shortest schedule is passed back to the simulation and the aircraft proceeds with the schedule.

The model of a single route is composed of several activities and resources. A diagram of a simple route is shown in Figure 4. Each box in the diagram represents a "stopping" activity. These activities require an appropriate amount of working MOG and have a fixed duration of the planning factor Ground Time (GT) plus a certain percentage. This percentage can be changed by the user and can affect how close the actual start and stop times are to the scheduled times. This will be examined later in this paper.

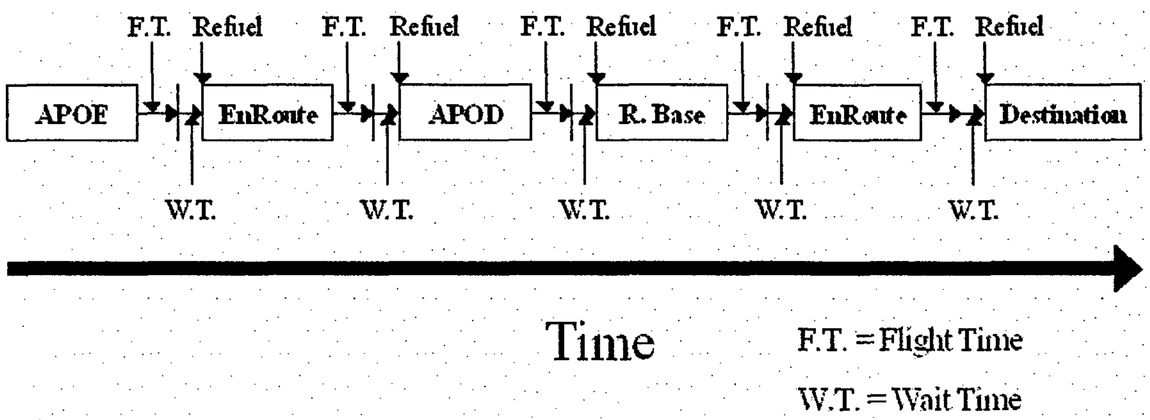

Figure 4: Regular Schedule Model

The Flight Times (F.T.) are activities representing the period of time needed to fly between locations that start at the time a stopping activity ends. The durations of these activities are calculated in the initialization block of the model using the aircraft block speed data and the distances between locations; they require no resources.

The Wait Times (W.T.) are periods of time of up to 30 minutes between the end of a Flight Time and the beginning of the following stopping activity. The search engine chooses the exact values of these times, which allow for some flexibility in creating a schedule. These Wait Times simulate the ability of an aircraft to wait in the air a short time to allow resources on the ground to become available.

Each stopping activity has a corresponding refueling activity with the same start time. The refueling activities require an appropriate amount of fuel resource from the corresponding airbase, which is calculated based on the flight time and aircraft fuel usage

data. The end time of the refueling activity is the end of the corresponding simulation day.

Each airbase has a corresponding fuel resource with maximum capacity of the maximum daily dispensing capacity. The fuel usage contained in the Expected Fuel Usage array is subtracted from this capacity over the corresponding simulation day using the maxCapacity function.

Modeling MOG

Modeling MOG is one of the most complicated aspects of this model. It is complicated by the fact that four different numbers and a logic type, not a simple number as is typical with modeling resources, describe MOG.

Stopping activities are modeled as requiring working MOG. Parking MOG is not considered in this model although the Wait Times do represent parking to some extent. Each location has two different types of resources related to MOG, NBWResource and WBWResource. For a location with a logic of "OR," the basic idea is that

$$\frac{\#NB\ Aircraft}{NBW\ MOG}+\frac{\#WB\ Aircraft}{WBW\ MOG}\leq 1.$$

To avoid using fractions, we multiply both sides by the denominators getting,

$$(\#NB\ Aircraft)*(WBW\ MOG)+(\#WB\ Aircraft)*(NBW\ MOG)\leq (NBW\ MOG)*(WBW\ MOG).$$

So, we define the NBWResource to have capacity (NBW MOG * WBW MOG) and the WBWResource have capacity 0 (i.e. WBWResource is not used). Each stopping activity of a wide body aircraft requires an amount of the NBWResource equal to the NBW

MOG, while each narrow body aircraft requires an amount of the NBWResource equal to the WBW MOG.

A location with a logic of "AND" has a NBWResource with capacity (NBW MOG + WBW MOG) and a WBWResource with capacity WBW MOG. Each narrow body aircraft requires one unit of the NBWResource. Each wide body aircraft requires one unit of the NBWResource and one unit of the WBWResource.

The maxCapacity function is used to model the presence of aircraft counted in the ExpectedOccupants array. These aircraft use the MOG resources as described above. As an example, consider a location with a working MOG of "5 AND 3". This location would have a NBWResource of capacity 8 and a WBWResource of capacity 3. If from the ExpectedOccupants array, 1 narrow body aircraft and 1 wide body aircraft expect to be at that location over a given block of time, maxCapacity would be used to set the available capacity of the NBWResource to 6 and the capacity of WBWResource to 2 over that block of time.

The Objective Function

In this example, if the name of the last activity were arriveAtDestination, the objective function would be:

minimize arriveAtDestination.end

(i.e., we are minimizing the total length of the schedule).

Solving this model takes less than $1/100^{th}$ of a second. After a schedule for each potential route has been created and the fastest one selected, the schedule, flight times, and cargo capacity are written to a flat-text data file. Control is then passed back to the simulation.

ASSIGNING CARGO TO AIRCRAFT

Cargo is assigned to aircraft by assigning TPFDD rows to aircraft in a sequence of iterations. The TPFDD row designates the APOE from which the cargo will be picked up and the corresponding APOD where the cargo will be delivered. The exact cargo on the assigned TPFDD row the aircraft will load depends on the type of aircraft and the amounts and types of cargo remaining on the TPFDD row when the aircraft arrives at the APOE.

The heuristic used to assign TPFDD rows to aircraft uses a strategy we call a "Rolling Event Horizon" method. Here, "event" refers to the pickup and delivery of cargo. Every three simulation hours, we answer the question "what is each aircraft going to do *next*?" More specifically, we answer the question "from which TPFDD row will each aircraft pickup cargo after it completes its current activity?" This activity may be the delivery of cargo at an APOD and the servicing at a Recovery Base, or the completion of maintenance at a Home Base.

The next TPFDD row assignments are stored in an array in the simulation, but are not permanently assigned to aircraft until the aircraft reach a Home Base, APOE, or Recovery Base. Not every assignment made in each iteration is made permanent.

In the first step of each iteration the simulation writes a flat-text data file containing the TPFDD minus the cargo that has been delivered and "claimed" and the "current positions" of all aircraft. Claimed cargo is cargo that has not been loaded or delivered, but which has an aircraft en route to load it. This notion of "claimed" cargo is used to prevent multiple aircraft from being assigned to a single load of cargo. The current positions data indicate the locations and approximate completion times of current activities of all the aircraft.

The process of assigning TPFDD rows is a basic four-step process:

1. Calculate the approximate times each aircraft could deliver the cargo on each TPFDD row based on the current position of the aircraft.
2. Define costs for assigning each TPFDD row to each aircraft.
3. Calculate the maximum number of aircraft of each type each row needs to deliver all remaining cargo on the row.
4. Solve an assignment problem.

Defining Costs

The method for defining costs is what we call a "ranking" strategy. The first step is to find the set of TPFDD rows, C_t, with non-trivial amounts of cargo remaining. The next step is to rank the rows in C_t according to their *RDD*'s. Ranks are integers ranging from 1 to $| C_t |$ and are denoted by *Rank*[*j*]. Two or more rows with the same *RDD* are given the same rank.

The next step is to have each row in C_t rank each aircraft within each aircraft type. Since C-5 and C-17 aircraft carry the same type of cargo, they are considered as one type. Aircraft are ranked according to two factors: approximate delivery day and distance from APOE when the current activity is complete. In general, the sooner the delivery day and distance, the lower the rank, and hence, the lower the cost.

The pseudo-code for defining the cost of assigning aircraft *i* to TPFDD row *j*, $\check{c}[i,j]$, is shown in Figure 5 where:

- V_1 = {tail numbers of WBP}
- V_2 = {tail numbers of WBC}
- V_3 = {tail numbers of C-5 and C-17}
- *DD*[*i,j*] = Delivery Day if aircraft *i* were to delivery cargo from row *j*
- *D*[*i,j*] = Distanee aircraft *i* would have to travel to pickup cargo from row *j*

The next step is to approximate the maximum number of each type of aircraft needed to deliver all remaining cargo in each row in C_t. Let:

- $WBP[j]$ = Maximum Number of WBP aircraft needed by TPFDD row j
- $WBC[j]$ = Maximum Number of WBC aircraft needed by TPFDD row j
- $C17[j]$ = Maximum Number of C-5 and C-17 aircraft needed by TPFDD row j
- MNA = Maximum total Number of All aircraft needed by all rows in C_t

If $MNA = 0$, all cargo has been delivered and the simulation is complete.

```
For m = 1 to 3:
   For j = 1 to |C_t|:
      D := V_m
      RankCounter := (Rank[j] – 1) * |V_m| + 1
      For k = 1 to |D|:
         Forall aircraft i ∈ D with minimum DD[i,j] ordered by increasing D[i,j]
            c[i,j] := RankCounter
            RankCounter := RankCounter + 1
            D := D \ {i}
         next i
      next k
   next j
next m
```

Figure 5: Pseudo-code for Defining Costs

The Assignment Model

The assignment model shown in Figure 6 has binary variables x_{ij} with $x_{ij} = 1$ meaning that aircraft i is assigned to TPFDD row j. Here:

- $CapWBC[j]$ = Average capacity of a WBC associated with row j

- $ABC[j]$ = Average Bulk Cargo carried by C-5's and C-17's associated with row j
- $Bulk[j]$ = Amount of Bulk cargo in row j
- $V = V_1 \cup V_2 \cup V_3$ (the set of all aircraft)

(Note: These first two parameters are technical details dealing with the way aircraft capacities are defined and the way aircraft are loaded.)

minimize

(2) $$\sum_i \sum_j x_{ij} c_{ij}$$

subject to

(3) $$\sum_{i \in V_1} x_{ij} \le MNP[j] \qquad \text{for all } j \in C_t$$

(4) $$\sum_{i \in V_2} x_{ij} \le MNC[j] \qquad \text{for all } j \in C_t$$

(5) $$\sum_{i \in V_3} x_{ij} \le C17[j] \qquad \text{for all } j \in C_t$$

(6) $$\sum_{i \in V_3} ABC[j] * x_{ij} + \sum_{i \in V_2} CapWBC[j] * x_{ij} \le CapWBC[j] + Bulk[j] \qquad \text{for all } j \in C_t$$

(7) $$\sum_{i \in V} \sum_{j \in C_t} x_{ij} = MNA$$

(8) $$\sum_{j \in C_t} x_{ij} \le 1 \qquad \text{forall } i \in V$$

(9) $$x_{ij} \in \{0,1\} \qquad \text{for all } i, j$$

Figure 6: The Assignment Model

The constraints ensure:

(3) - (5) No row is assigned too many aircraft of any type.

(6) The approximate total bulk cargo carrying capacity of all aircraft assigned to a row can't be in excess of the amount of bulk cargo available by more than the capacity of a WBC.

(7) The maximum possible number of assignments is made.

(8) Each aircraft is assigned to at most one TPFDD row.

(9) The variables are binary.

Because of the bulk carrying capacity of C-5's and C-17's and this interaction with the number of WBC aircraft chosen, constraint (7) may make the model infeasible. In the event this happens, *MNA* is decreased by one and the model is re-solved. This procedure is repeated until the model is feasible. At the extreme, setting *MNA* = 0 guarantees a solution of all variables equal to 0.

RESULTS

The simulation is complete when all cargo has been delivered. Each iteration of the assignment heuristic takes less than a second to complete. Overall, with the TPFDD provided to us by AMC, the simulation lasts about 1100 simulated hours in approximately 12 minutes real time.

The overall results in terms of the delivery of cargo are shown in Figure 7. The curve labeled "TPFDD Data" shows the cumulative tons of cargo required to be delivered by each simulation day in the TPFDD. The curve labeled "Heuristic" shows the cumulative tons of cargo actually delivered each day in the simulation. Ideally, the "Heuristic" curve would be above the "TPFDD" curve meaning that the transportation demands were met. However, because of the limited number of C-5 and C-17 aircraft, we were unable to meet demands.

To compare this assignment heuristic to a much simpler one, we devised a "look-down list" strategy that has been used by analysts at AMC in other simulations. In this strategy,

the rows of the TPFDD are arranged in ascending order according to their *RDD*'s. When an aircraft needs its next TPFDD row assignment, it simply looks down the TPFDD until it finds a row with some remaining cargo.

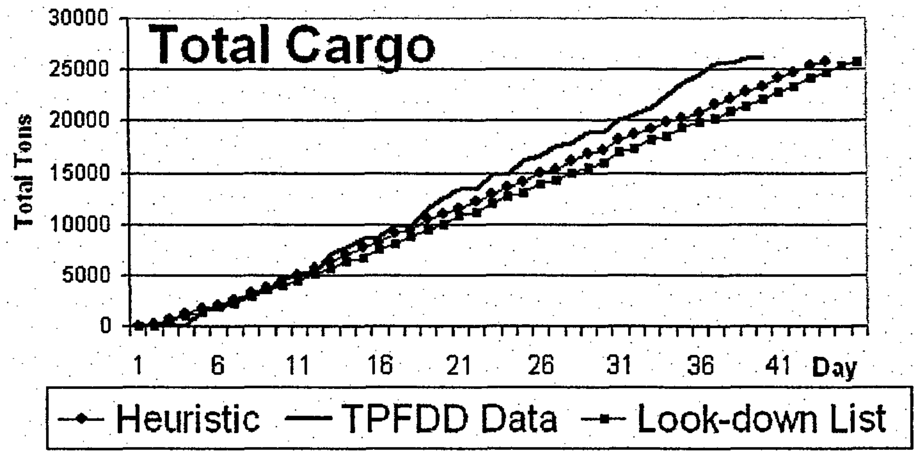

Figure 7: Cargo Delivery Results

Figure 7 shows a comparison of the results from the heuristic to the look-down list strategy. As can be seen, the heuristic strategy yielded significantly better results than the look-down list strategy. From the slopes of the linear best-fit lines, we see that using the heuristic yielded an average delivery of 572 tons per day while the look-down list strategy yielded 539 tons per day, about 6% lower.

Figures 8 and 9 show additional comparisons between the two strategies regarding the late deliveries. Figure 8 shows the average lateness in terms of days for all late deliveries. This graphs shows that the heuristic resulted in late deliveries that were less late on average than the look-down list strategy. Figure 9 shows the percent of all deliveries that were made late. Again, we see that the heuristic strategy produced a lower percentage of late deliveries than the look-down list strategy.

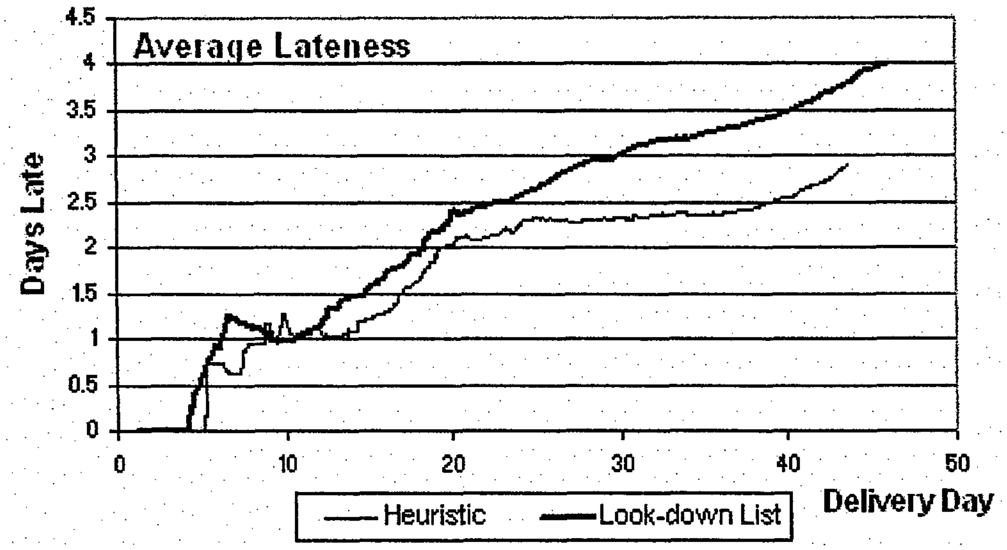

Figure 8: Comparison of Lateness

Quality of Schedules

When an aircraft has completed a schedule, the columns of the schedule labeled expected start and end times contain the actual times each activity started and began. The scheduled times did not change. The schedule is exported to a flat-text data file from where it can be analyzed in Excel after the simulation.

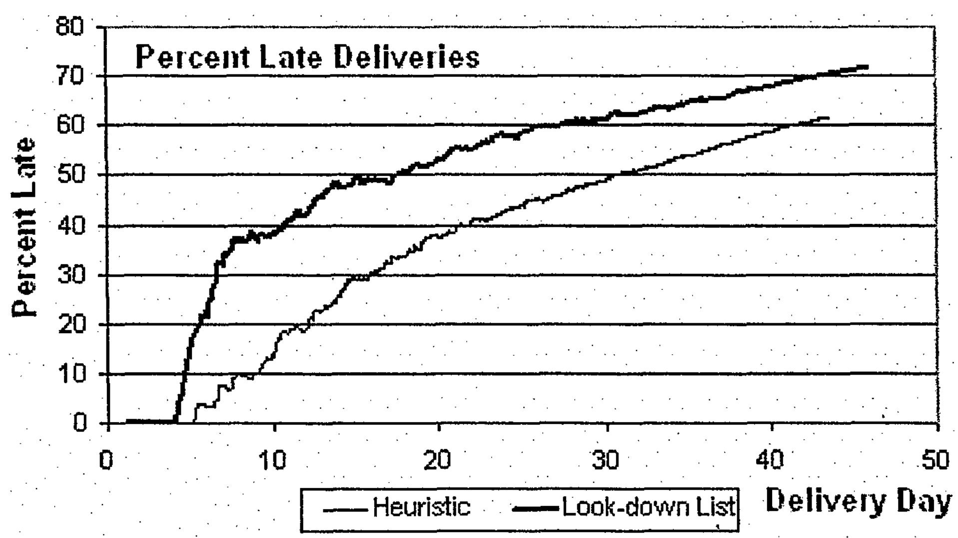

Figure 9: Comparison of Late Deliveries

Of particular interest is how close the actual start times are to the schedule times. Ideally these times would be very close together. In creating the schedules, the durations of the stopping activities were set to the planning factor ground time (GT) plus a percentage. This percentage can be set by the user and has a great effect on the quality of the schedules.

To measure the quality of the schedules we calculated the distribution of the differences of the actual and scheduled start times. A negative (positive) difference means that the activity started ahead of (behind) schedule.

Figure 10 shows the distributions for different values of the extra percentage. Using 5% extra time is equivalent to using the mean ground time in the simulation. As can be seen, the distributions using 0% or 5% are skewed to the right, indicating that activities tended to start late. The distribution for 15% is skewed left, indicating that activities tended to start early. The distribution for 10% shows a spike at 0%, as is expected. Removing this

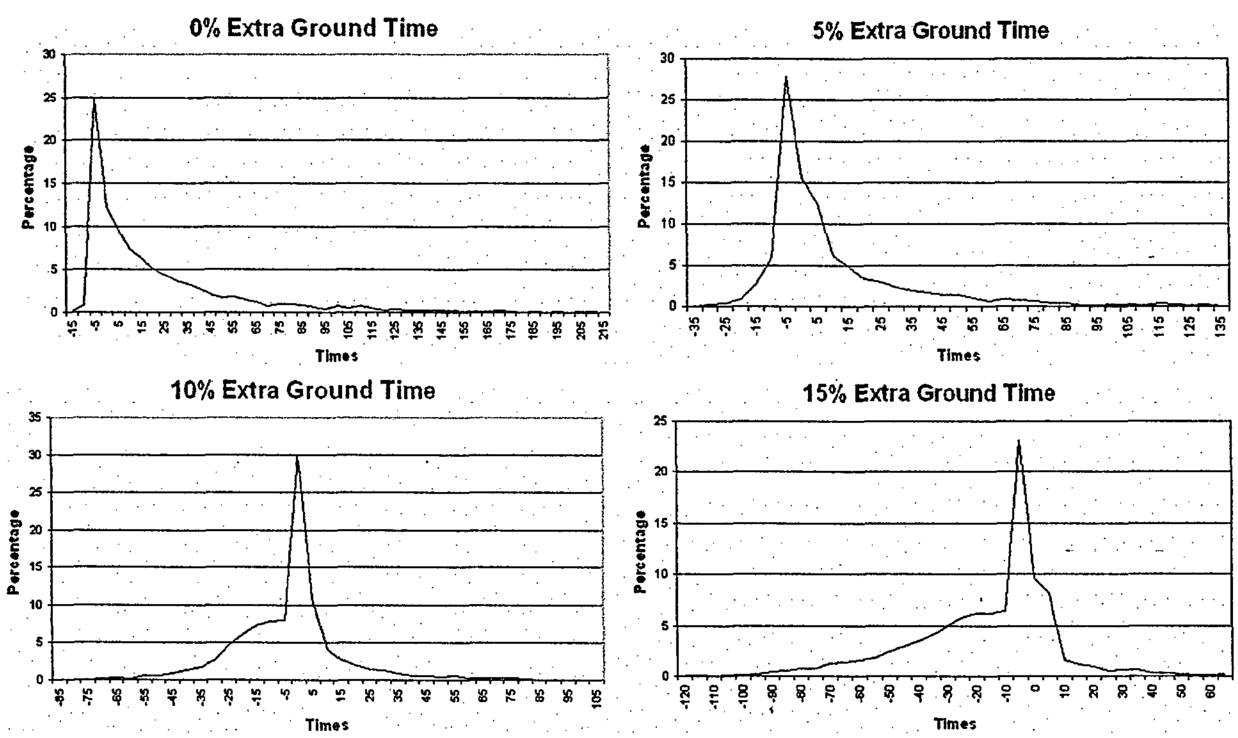

Figure 10: Distributions of Differences in Actual and Scheduled Start Times

spike and "smoothing" out the curve shows the distribution is relatively normal, centered near 0, and the majority of the differences fall between plus or minus 35 minutes. This indicates that most activities tended to start within about a half hour of their scheduled times, which shows a high quality of schedules.

CONCLUSIONS

In conclusion we have shown how optimization techniques can be embedded in the framework of a stochastic discrete-event simulation of a complex military airlift network using off-the-shelf technology. We have shown how the very complicated decisions of assigning cargo to aircraft and choosing routes can be modeled using linear and non-linear models. We have also shown how these methods produce considerably better results than the commonly used look-down list methods for making decisions.

We hope that our results will convince military transportation analysts of the benefits of using optimization in their analysis.

REFERENCES

Cusick, Travis W. 2000. Simulation and Optimization of a Military Airfield System. Ph.D. Dissertation. Department of Systems Science and Mathematics, Washington University in St. Louis.

Grindey, Gregory J. 2002. Intelligent Transportation Scheduling: Heuristic and Sequential Optimization of Simulated Transportation Systems. Ph.D. Dissertation. Department of Systems Science and Mathematics, Washington University in St. Louis.

Morton, D.P., R.E. Rosenthal, and L.T. Weng. 1996. Optimization Modeling for Airlift Mobility, Military Operations Research 1, Vol 1, No 4, 49-67.

Powell, Warren B. 2001. The Optimizing-Simulator: An Analysis Technology for Dynamic Resource Management. Unpublished technical paper. CASTLE Laboratory, Princeton University.

Rink, Katherine A. 1998. Adaptation of Shortest Path Algorithms to Mobility Problems. Ph.D. Disseration. Department of Systems Science and Mathematics, Washington University in St. Louis.